arXiv:1611.01708v1 [stat.ML] 5 Nov 2016

Detecting Dependencies in High-Dimensional, Sparse Databases Using Probabilistic Programming and Non-parametric Bayes

Feras Saad
Computer Science & AI Laboratory
Massachusetts Institute of Technology

Vikash Mansinghka
Department of Brain & Cognitive Sciences
Massachusetts Institute of Technology

Abstract

Sparse databases with hundreds of variables are commonplace. In this setting, it is both statistically and computationally challenging to detect true predictive relationships between variables and also to suppress false positives. This paper proposes a new approach to dependency detection that combines probabilistic programming, information theory, and non-parametric Bayesian modeling. The key ideas are to (i) build an ensemble of joint probability models for the whole database via approximate posterior inference in CrossCat, a non-parametric factorial mixture; (ii) identify independencies by analyzing model structures; and (iii) report the distribution on conditional mutual information induced by posterior uncertainty over the ensemble of models. This paper presents experiments showing that the approach finds relationships that pairwise correlation misses, including context-specific independencies, on databases of mathematics exam scores and global indicators of macroeconomic development.

1 Introduction

Sparse databases with hundreds of variables are commonplace. In these settings, it can be both statistically and computationally challenging to detect predictive relationships between variables [4]. First, the data may be incomplete, and obtained from multiple primary sources, often of non-uniform quality. Cleaning and imputation are often needed before pairwise statistics can be calculated. Second, parametric modeling assumptions that underlie standard hypothesis testing techniques may not be appropriate; relationships may be nonlinear, multivariate, and/or heteroskedastic, and thus missed by pairwise correlations. Third, as the number of variables grows, it becomes harder to detect true relationships without also finding false positives. Many frequentist approaches have been proposed (see [16, Table 1] for a summary), but they do not produce probabilistically coherent measures of uncertainty, and can struggle with data sparsity. This paper proposes a new approach to dependence detection that combines probabilistic programming, information theory, and non-parametric Bayesian modeling. The approach is summarized in Figure 1. Queries about the conditional mutual information (CMI) between variables of interest are expressed using the Bayesian Query Language [17], an SQL-like probabilistic programming language. Approximate inference with CrossCat [18] is used to produce an ensemble of joint probability models, which are analyzed for structural (in)dependencies. For model structures in which dependence cannot be ruled out, the CMI is estimated via Monte Carlo. The method is scalable to high-dimensional data; it can be used for exploratory analysis without requiring expensive CMI estimation for all pairs of variables. It applies to heterogeneously typed, incomplete datasets with minimal pre-processing [18].

This paper shows empirically that the proposed approach is effective on real-world databases with hundreds of variables and a missing data rate of $\sim$35%, detecting common-sense predictive relationships that are missed by pairwise correlation while suppressing spurious relationships that other methods purport to detect. This paper also shows how to characterize the strength of context-specific (in)dependencies on real-world databases of mathematics exam scores and of global macroeconomic indicators.

2 Drawing Bayesian inferences about conditional mutual information

Let $\boldsymbol{x} = (x_1, x_2, \ldots, x_D)$ denote an absolutely-continuous D-dimensional random vector, whose subvectors we denote $\boldsymbol{x}_{\mathcal{A}} = \{x_i : i \in \mathcal{A}\}$ with joint probability density $p_{\mathcal{G}}(\boldsymbol{x}_{\mathcal{A}})$. The symbol $\mathcal{G}$ refers to an arbitrary specification for the "generative" process of

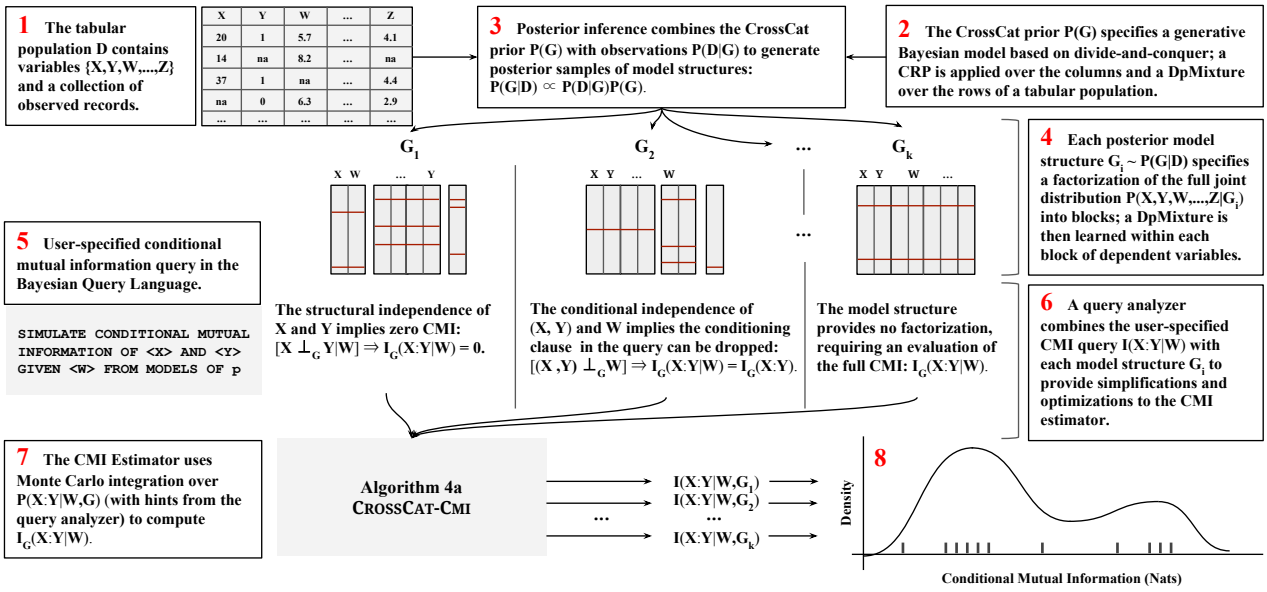
1 The tabular population D contains variables {X,Y,W,....,Z} and a collection of observed records.
X Y W ... Z
20 1 5.7 ... 4.1
14 na 8.2 ... na
37 1 na ... 4.4
na 0 6.3 ... 2.9
... ...
3 Posterior inference combines the CrossCat prior P(G) with observations P(D|G) to generate posterior samples of model structures: P(G|D) ∝ P(D|G)P(G).
2 The CrossCat prior P(G) specifies a generative Bayesian model based on divide-and-conquer; a CRP is applied over the columns and a DpMixture over the rows of a tabular population.
G1
G2
...
Gk
4 Each posterior model structure Gi ~ P(G|D) specifies a factorization of the full joint distribution P(X,Y,W,....,Z|Gi) into blocks; a DpMixture is then learned within each block of dependent variables.
5 User-specified conditional mutual information query in the Bayesian Query Language.
SIMULATE CONDITIONAL MUTUAL INFORMATION OF <X> AND <Y> GIVEN <W> FROM MODELS OF p
The structural independence of X and Y implies zero CMI: [X ⊥G Y|W] ⇒ IG(X:Y|W) = 0.
The conditional independence of (X, Y) and W implies the conditioning clause in the query can be dropped: [(X ,Y) ⊥G W] ⇒ IG(X:Y|W) = IG(X:Y).
The model structure provides no factorization, requiring an evaluation of the full CMI: IG(X:Y|W).
6 A query analyzer combines the user-specified CMI query I(X:Y|W) with each model structure Gi to provide simplifications and optimizations to the CMI estimator.
7 The CMI Estimator uses Monte Carlo integration over P(X:Y|W,G) (with hints from the query analyzer) to compute IG(X:Y|W).
Algorithm 4a CrossCat-Cmi
I(X:Y|W,G1)
I(X:Y|W,G2)
...
I(X:Y|W,Gk)
8
Density
Conditional Mutual Information (Nats)

Figure 1: Workflow for computing posterior distributions of the conditional mutual information of arbitrary variables in a data table. Details outlining the integration of this workflow into BayesDB, a probabilistic programming platform for data analysis, and new syntaxes in the Bayesian Query Language (block 5) for simulating from the distribution over CMI, are provided in the supplement.

$\boldsymbol{x}$, and parameterizes all its joint and conditional distributions. The *mutual information* (MI) of $\boldsymbol{x}_\mathcal{A}$ and $\boldsymbol{x}_\mathcal{B}$ (under $\mathcal{G}$) is defined in the usual way [5]:

$$\mathcal{I}_\mathcal{G}(\boldsymbol{x}_\mathcal{A}{:}\boldsymbol{x}_\mathcal{B}) = \mathop{\mathbb{E}}_{(\boldsymbol{x}_\mathcal{A},\boldsymbol{x}_\mathcal{B})}\left[\log\left(\frac{p_\mathcal{G}(\boldsymbol{x}_\mathcal{A},\boldsymbol{x}_\mathcal{B})}{p_\mathcal{G}(\boldsymbol{x}_\mathcal{A})p_\mathcal{G}(\boldsymbol{x}_\mathcal{B})}\right)\right]. \quad (1)$$

This quantity can be interpreted as the KL-divergence between the joint distribution $p_\mathcal{G}(\boldsymbol{x}_\mathcal{A},\boldsymbol{x}_\mathcal{B})$ and product of marginals $p_\mathcal{G}(\boldsymbol{x}_\mathcal{A})p_\mathcal{G}(\boldsymbol{x}_\mathcal{B})$, and is a well-established measure for both the existence and strength of dependence between $\boldsymbol{x}_\mathcal{A}$ and $\boldsymbol{x}_\mathcal{B}$ (Section 2.2). Given an observation of the variables $\{\boldsymbol{x}_\mathcal{C}{=}\hat{\boldsymbol{x}}_\mathcal{C}\}$, the *conditional mutual information* (CMI) of $\boldsymbol{x}_\mathcal{A}$ and $\boldsymbol{x}_\mathcal{B}$ given $\{\boldsymbol{x}_\mathcal{C}{=}\hat{\boldsymbol{x}}_\mathcal{C}\}$ is defined analogously:

$$\mathcal{I}_\mathcal{G}(\boldsymbol{x}_\mathcal{A}{:}\boldsymbol{x}_\mathcal{B}|\boldsymbol{x}_\mathcal{C}{=}\hat{\boldsymbol{x}}_\mathcal{C}) = \mathop{\mathbb{E}}_{(\boldsymbol{x}_\mathcal{A},\boldsymbol{x}_\mathcal{B})|\hat{\boldsymbol{x}}_\mathcal{C}}\left[\log\left(\frac{p_\mathcal{G}(\boldsymbol{x}_\mathcal{A},\boldsymbol{x}_\mathcal{B}|\hat{\boldsymbol{x}}_\mathcal{C})}{p_\mathcal{G}(\boldsymbol{x}_\mathcal{A}|\hat{\boldsymbol{x}}_\mathcal{C})p_\mathcal{G}(\boldsymbol{x}_\mathcal{B}|\hat{\boldsymbol{x}}_\mathcal{C})}\right)\right]. \quad (2)$$

Estimating the mutual information between the variables of $\boldsymbol{x}$ given a dataset $\mathcal{D} = \{\hat{\boldsymbol{x}}_i\}_{i=1}^N$ is an open problem in the literature. Various parametric and non-parametric methods for estimating MI exist [20, 21, 14]; see [23] for a comprehensive review. Traditional approaches typically construct a point estimate $\hat{\mathcal{I}}(\boldsymbol{x}_\mathcal{A}{:}\boldsymbol{x}_\mathcal{B})$ (and confidence intervals) around a "true value" of $\mathcal{I}(\boldsymbol{x}_\mathcal{A}{:}\boldsymbol{x}_\mathcal{B})$. In this paper, we instead take a non-parametric Bayesian approach, where the mutual information itself is a derived random variable; a similar interpretation was recently developed in independent work by [15]. The randomness of mutual information arises from treating the data generating process and parameters $\mathcal{G}$ as a random variable, whose prior distribution we denote π. Composing $\mathcal{G}$ with the measurable function $h : \hat{\mathcal{G}} \mapsto \mathcal{I}_{\hat{\mathcal{G}}}(\boldsymbol{x}_\mathcal{A}{:}\boldsymbol{x}_\mathcal{B})$ induces the random variable $h(\mathcal{G}) \equiv \mathcal{I}_\mathcal{G}(\boldsymbol{x}_\mathcal{A}{:}\boldsymbol{x}_\mathcal{B})$. The distribution of the MI is thus an expectation under measure π:

$$\begin{aligned}\mathbb{P}\left[\mathcal{I}_\mathcal{G}(\boldsymbol{x}_\mathcal{A}{:}\boldsymbol{x}_\mathcal{B}) \in S\right] &= \int \mathbb{I}\left[\mathcal{I}_{\hat{\mathcal{G}}}(\boldsymbol{x}_\mathcal{A}{:}\boldsymbol{x}_\mathcal{B}) \in S\right]\pi(d\hat{\mathcal{G}}) \\ &= \mathop{\mathbb{E}}_{\hat{\mathcal{G}}\sim\pi}\left[\mathbb{I}\left[\mathcal{I}_{\hat{\mathcal{G}}}(\boldsymbol{x}_\mathcal{A}{:}\boldsymbol{x}_\mathcal{B}) \in S\right]\right], \quad (3)\end{aligned}$$

for all Borel sets S. Given a dataset $\mathcal{D}$, the posterior distribution of the mutual information, $\mathbb{P}(\mathcal{I}_\mathcal{G}(\boldsymbol{x}_\mathcal{A}{:}\boldsymbol{x}_\mathcal{B}) \in S|\mathcal{D})$, is constructed by using the posterior $\pi(\cdot|\mathcal{D})$ instead of π in (3). The distribution over conditional mutual information (2) is computed in an analogous manner, where the CMI is used inside the expectation.

2.1 Estimating conditional mutual information with generative population models

It is straightforward to produce Monte Carlo estimates of CMI for all *generative population models*, a probabilistic programming formalism for characterizing the data generating process of an infinite array of realizations of the random vector $\boldsymbol{x} = (x_1, x_2, \ldots, x_D)$ introduced in [17, 27]. Listing 1 summarizes the key elements of the GPM interface.

Listing 1 GPM interface for simulating from and assessing the density of conditional and marginal distributions of a random vector $\boldsymbol{x}$.

SIMULATE($\mathcal{G}$, query: $\mathcal{Q}=\{q_j\}$, condition: $\hat{\boldsymbol{x}}_{\mathcal{E}}=\{\hat{x}_{e_j}\}$)
Return a sample $\mathbf{s} \sim p_{\mathcal{G}}(\boldsymbol{x}_{\mathcal{Q}}|\hat{\boldsymbol{x}}_{\mathcal{E}}, \mathcal{D})$.

LOGPDF($\mathcal{G}$, query: $\hat{\boldsymbol{x}}_{\mathcal{Q}}=\{\hat{x}_{q_j}\}$, condition: $\hat{\boldsymbol{x}}_{\mathcal{E}}=\{\hat{x}_{e_j}\}$)
Return the joint log density $p_{\mathcal{G}}(\hat{\boldsymbol{x}}_{\mathcal{Q}}|\hat{\boldsymbol{x}}_{\mathcal{E}}, \mathcal{D})$

These two interface procedures can be combined to derive a simple Monte Carlo estimator for the CMI (2), shown in Algorithm 2a.

Algorithm 2a GPM-CMI

Require: GPM $\mathcal{G}$; query $\mathcal{A}$, $\mathcal{B}$; condition $\hat{\boldsymbol{x}}_{\mathcal{C}}$; accuracy T
Ensure: Monte Carlo estimate of $\mathcal{I}_{\mathcal{G}}(\boldsymbol{x}_{\mathcal{A}}:\boldsymbol{x}_{\mathcal{B}}|\boldsymbol{x}_{\mathcal{C}}=\hat{\boldsymbol{x}}_{\mathcal{C}})$
1: **for** $t = 1, \ldots, T$ **do**
2: $\quad(\hat{\boldsymbol{x}}_{\mathcal{A}}, \hat{\boldsymbol{x}}_{\mathcal{B}}) \leftarrow \text{SIMULATE}(\mathcal{G}, \mathcal{A} \cup \mathcal{B}, \hat{\boldsymbol{x}}_{\mathcal{C}})$
3: $\quad m^t_{\mathcal{A}\cup\mathcal{B}} \leftarrow \text{LOGPDF}(\mathcal{G}, \hat{\boldsymbol{x}}_{\mathcal{A}\cup\mathcal{B}}, \hat{\boldsymbol{x}}_{\mathcal{C}})$
4: $\quad m^t_{\mathcal{A}} \leftarrow \text{LOGPDF}(\hat{\boldsymbol{x}}_{\mathcal{A}}, \hat{\boldsymbol{x}}_{\mathcal{C}})$
5: $\quad m^t_{\mathcal{B}} \leftarrow \text{LOGPDF}(\hat{\boldsymbol{x}}_{\mathcal{B}}, \hat{\boldsymbol{x}}_{\mathcal{C}})$
6: **return** $\frac{1}{T}\sum_{t=1}^{T}\left(m^t_{\mathcal{A}\cup\mathcal{B}} - (m^t_{\mathcal{A}} + m^t_{\mathcal{B}})\right)$

While GPM-CMI is an unbiased and consistent estimator applicable to any probabilistic model implemented as a GPM, its quality in detecting dependencies is tied to the ability of $\mathcal{G}$ to capture patterns from the dataset $\mathcal{D}$; this paper uses baseline non-parametric GPMs built using CrossCat (Section 3).

2.2 Extracting conditional independence relationships

The (conditional) mutual information can be used to give straightforward definitions for marginal, context-specific, and conditional independence:

Marginal Independence It is straightforward to see that $(\boldsymbol{x}_{\mathcal{A}} \perp\!\!\!\perp_{\mathcal{G}} \boldsymbol{x}_{\mathcal{B}})$ if and only if $\mathcal{I}_{\mathcal{G}}(\boldsymbol{x}_{\mathcal{A}}:\boldsymbol{x}_{\mathcal{B}}) = 0$.

Context-Specific Independence If the event $\{\boldsymbol{x}_{\mathcal{C}}=\hat{\boldsymbol{x}}_{\mathcal{C}}\}$ decouples $\boldsymbol{x}_{\mathcal{A}}$ and $\boldsymbol{x}_{\mathcal{B}}$ completely, then they are said to be independent "in the context" of $\hat{\boldsymbol{x}}_{\mathcal{C}}$, written $(\boldsymbol{x}_{\mathcal{A}} \perp\!\!\!\perp_{\mathcal{G}} \boldsymbol{x}_{\mathcal{B}}|\{\boldsymbol{x}_{\mathcal{C}}=\hat{\boldsymbol{x}}_{\mathcal{C}}\})$ [3]. This condition is equivalent to the CMI from (2) equaling zero. Thus by estimating CMI, it is possible to detect finer-grained independencies than can be detected by analyzing the graph structure of a learned Bayesian network [29].

Conditional Independence If context-specific independence holds for all possible observation sets $\{\boldsymbol{x}_{\mathcal{C}}=\hat{\boldsymbol{x}}_{\mathcal{C}}\}$, then $\boldsymbol{x}_{\mathcal{A}}$ and $\boldsymbol{x}_{\mathcal{B}}$ are *conditionally independent* given $\boldsymbol{x}_{\mathcal{C}}$, written $(\boldsymbol{x}_{\mathcal{A}} \perp\!\!\!\perp_{\mathcal{G}} \boldsymbol{x}_{\mathcal{B}}|\boldsymbol{x}_{\mathcal{C}})$. By the non-negativity of CMI, conditional independence is equivalent to the CMI of $\boldsymbol{x}_{\mathcal{A}}$ and $\boldsymbol{x}_{\mathcal{B}}$ (marginalizing out $\boldsymbol{x}_{\mathcal{C}}$)

$$\mathcal{I}_{\mathcal{G}}(\boldsymbol{x}_{\mathcal{A}}:\boldsymbol{x}_{\mathcal{B}}|\boldsymbol{x}_{\mathcal{C}}) = \mathop{\mathbb{E}}_{\hat{\boldsymbol{x}}_{\mathcal{C}}}\left[\mathcal{I}_{\mathcal{G}}(\boldsymbol{x}_{\mathcal{A}}:\boldsymbol{x}_{\mathcal{B}}|\boldsymbol{x}_{\mathcal{C}} = \hat{\boldsymbol{x}}_{\mathcal{C}})\right], \qquad (4)$$

equaling zero. In Section 4, we compare posterior CMI curves for both context-specific and conditional queries, illustrating the difference between the two.

3 Building generative population models for CMI estimation with non-parametric Bayes

Our approach to estimating the CMI requires a prior π and model class $\mathcal{G}$ which is flexible enough to emulate an arbitrary joint distribution over $\boldsymbol{x} = (x_1, \ldots, x_D)$, and tractable enough to efficiently implement Algorithm 2a for arbitrary sub-vectors of $\boldsymbol{x}$. First, consider the special case of a multivariate Dirichlet process mixture model [11], here abbreviated DpMixture. Letting L_d denote the likelihood function for variable x_d, V_d a prior over the parameters of L_d, and $\boldsymbol{\lambda}_d$ the hyperparameters of V_d, we summarize the prior generative process for data $\mathcal{D} = \{\boldsymbol{x}_i\}_{i=1}^N$ by:

DPMIXTURE-PRIOR

$$\alpha \sim \text{GAMMA}(1, 1)$$
$$\boldsymbol{z} = (z_1, \ldots, z_N) \sim \text{CRP}(\cdot|\alpha)$$
$$\boldsymbol{\phi}_{[d,k]} \sim V_d(\cdot|\boldsymbol{\lambda}_d) \qquad d \in [D], k \in \text{UNIQUE}(\boldsymbol{z})$$
$$x_{[i,d]} \sim L_d(\cdot|\boldsymbol{\phi}_{[d,z_i]}) \qquad i \in [N], d \in [D]$$

We refer to [7, 13] for algorithms for posterior inference, and assume we have a posterior sample $\hat{\mathcal{G}} = (\alpha, \boldsymbol{z}_{[1:N]}, \{\boldsymbol{\phi}_d\})$ of all latent parameters in the DpMixture given $\mathcal{D}$. To compute the CMI of an arbitrary query pattern $\mathcal{I}_{\hat{\mathcal{G}}}(\boldsymbol{x}_{\mathcal{A}} : \boldsymbol{x}_{\mathcal{B}}|\boldsymbol{x}_{\mathcal{C}} = \hat{\boldsymbol{x}}_{\mathcal{C}})$ using Algorithm 2a, we need implementations of SIMULATE and LOGPDF for DpMixture $\hat{\mathcal{G}}$. These two procedures are summarized in Algorithms 3a and 3b.

Algorithm 3a DPMIXTURE-SIMULATE

Require: DpMixture $\mathcal{G}$; target $\mathcal{A}$; condition $\hat{\boldsymbol{x}}_{\mathcal{C}}$
Ensure: joint sample $\hat{\boldsymbol{x}}_{\mathcal{A}} \sim p_{\mathcal{G}}(\cdot|\hat{\boldsymbol{x}}_{\mathcal{C}})$
1: $(l_i)_{i=1}^{K+1} \leftarrow \text{DPMIXTURE-CLUSTER-POSTERIOR}(\mathcal{G}, \hat{\boldsymbol{x}}_{\mathcal{C}})$
2: $z_{N+1} \sim \text{CATEGORICAL}(l_1, \ldots, l_{K+1})$
3: **for** $a \in \mathcal{A}$ **do**
4: $\quad\hat{x}_a \sim L_a(\cdot|\boldsymbol{\phi}_{[a,z_{N+1}]})$
5: **return** $\hat{\boldsymbol{x}}_{\mathcal{A}}$

Algorithm 3b DPMIXTURE-LOGPDF

Require: DpMixture $\mathcal{G}$; target $\hat{\boldsymbol{x}}_{\mathcal{A}}$; condition $\hat{\boldsymbol{x}}_{\mathcal{C}}$
Ensure: log density $p_{\mathcal{G}}(\hat{\boldsymbol{x}}_{\mathcal{A}}|\hat{\boldsymbol{x}}_{\mathcal{C}})$
1: $(l_i)_{i=1}^{K+1} \leftarrow \text{DPMIXTURE-CLUSTER-POSTERIOR}(\mathcal{G}, \hat{\boldsymbol{x}}_{\mathcal{C}})$
2: **for** $k = 1, \ldots, K+1$ **do**
3: $\quad t_k \leftarrow \prod_{a\in\mathcal{A}} L_a(\hat{x}_a|\boldsymbol{\phi}_{[a,k]})$
4: **return** $\log\left(\sum_{k=1}^{K+1}(t_k l_k)\right)$

Algorithm 3c DPMIXTURE-CLUSTER-POSTERIOR

Require: DpMixture $\mathcal{G}$; condition $\hat{\boldsymbol{x}}_{\mathcal{C}}$;
Ensure: $\{p_{\mathcal{G}}(z_{N+1} = k) : 1 \le k \le \max(\boldsymbol{z}_{1:N}) + 1\}$
1: $K \leftarrow \max(\boldsymbol{z}_{1:N})$
2: **for** $k = 1, \ldots, K+1$ **do**
3: $\quad n_k \leftarrow \begin{cases} |\{\boldsymbol{x}_i \in \mathcal{D} : z_i = k\}| & \text{if } k \le K \\ \alpha & \text{if } k = K+1 \end{cases}$
4: $\quad l_k \leftarrow \left(\prod_{c\in\mathcal{C}} L_c(\hat{x}_c|\boldsymbol{\phi}_{[c,k]})\right) n_k$
5: **return** $(l_1, \ldots, l_{K+1}) / \sum_{k=1}^{K+1}(l_k)$

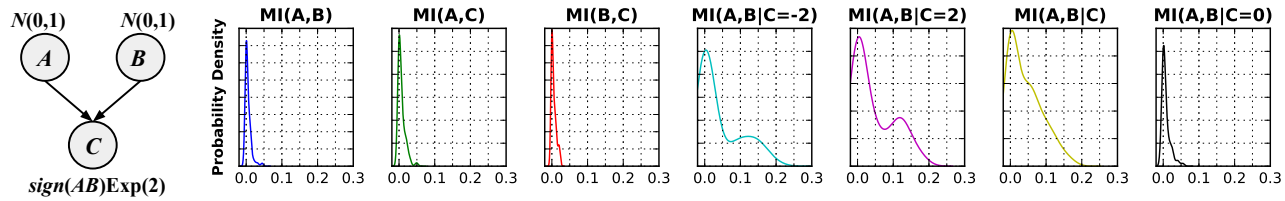

(a) Ground truth "common-effect" generator [28].

(b) The first three plots verify that A, B, and C are marginally independent. The next three plots show that conditioning on C "couples" the parents A and B (both for fixed values of $C \in \{2, -2\}$, and marginalizing over all C). The last plot shows that $\{C = 0\}$ does not couple A and B, due to symmetry of signum.

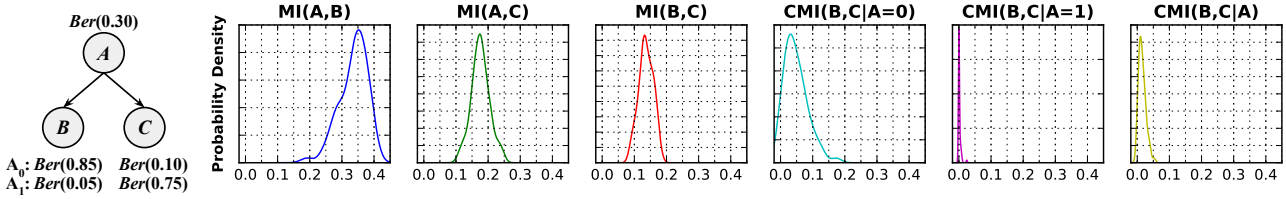

(c) Ground truth "common-cause" generator.

(d) The first three plots verify that A, B, and C are marginally dependent. The next three plots show that conditioning on A "decouples" the children B and C; the decoupling is weaker for $\{A = 0\}$, because it is 3.4 nats less likely that $\{A = 1\}$. The final plot shows the weighted CMI under these two possibilities.

Figure 2: Posterior distributions of CMI under the CrossCat posterior, given 100 data points from canonical Bayes net structures. Distributions peaked at 0 indicate high probability of (conditional) independence. In both cases, the posterior CMI correctly detects the marginal and conditional independence structures in the "ground truth" Bayes nets, despite the fact that both "common-cause" and "common-effect" structures are not in the (structural) hypothesis space of a DpMixture.

The subroutine DpMixture-Cluster-Posterior is used for sampling (in DpMixture-Simulate) and marginalizing over (in DpMixture-LogPdf) the nonparametric mixture components. Moreover, if L_d and V_d form a conjugate likelihood-prior pair, then invocations of $L_d(\hat{x}_d|\boldsymbol{\phi}_{[d,k]})$ in Algorithms 3a:4 and 3b:3 can be Rao-Blackwellized by conditioning on the sufficient statistics of data in cluster k, thus marginalizing out $\boldsymbol{\phi}_{[d,k]}$ [25]. This optimization is important in practice, since analytical marginalization can be obtained in closed-form for several likelihoods in the exponential family [8]. Finally, to approximate the posterior distribution over CMI in (4), it suffices to recompute DpMixture-Cmi using a set of posterior samples $\left\{\hat{\mathcal{G}}_1, \ldots, \hat{\mathcal{G}}_H\right\} \sim^{\text{iid}} \pi(\cdot|\mathcal{D})$. Figure 2 shows various posterior CMI distributions from the DpMixture successfully recovering the marginal and conditional independencies in two canonical Bayesian networks.

3.1 Inducing Sparse Dependencies Using The CrossCat Prior

A single DpMixture assumes that all variables $\boldsymbol{x} = (x_1, \ldots, x_D)$ are (structurally) marginally dependent, where their joint distribution factorizes conditioned on the mixture assignment z. In high-dimensional datasets, imposing full structural dependence among all the x_j is often a conservative assumption. Moreover, while the Monte Carlo error of Algorithm 2a does not scale with the dimensionality D, its runtime scales linearly for the DpMixture, and so estimating the CMI is likely to be prohibitively expensive. We relax this constraint by using the CrossCat prior [18], a structure learning meta-GPM which induces sparsity over the structural dependencies in the variables of $\boldsymbol{x}$. In particular, CrossCat posits a factorization of $\boldsymbol{x}$ according to a *variable* partition $\boldsymbol{\gamma} = \left\{\mathcal{V}_1, \ldots, \mathcal{V}_{|\boldsymbol{\gamma}|}\right\}$, where $\mathcal{V}_i \subseteq [D]$. For all $i \neq j$, all variables in block $\mathcal{V}_i$ are mutually (marginally and conditionally) independent of all variables in $\mathcal{V}_j$. Thus factorization of $\boldsymbol{x}$ is:

$$p_{\mathcal{G}}(\boldsymbol{x}|\mathcal{D}) = \prod_{\mathcal{V} \in \boldsymbol{\gamma}} p_{\mathcal{G}_{\mathcal{V}}}(\boldsymbol{x}_{\mathcal{V}}|\mathcal{D}_{\mathcal{V}}). \tag{5}$$

Within block $\mathcal{V}$, the variables $\boldsymbol{x}_{\mathcal{V}} = \{x_d : d \in \mathcal{V}\}$ are distributed according to a vanilla DpMixture; subscripts with $\mathcal{V}$ (such as $\mathcal{G}_{\mathcal{V}}$) now index a set of block-specific DpMixture parameters. The joint (predictive) density $p_{\mathcal{G}_{\mathcal{V}}}$ is given by Algorithm 3b:

$$p_{\mathcal{G}_{\mathcal{V}}}(\boldsymbol{x}_{\mathcal{V}}|\mathcal{D}) = \sum_{k=1}^{K_{\mathcal{V}}+1} \left(\frac{n_{[\mathcal{V},k]} \prod_{d \in \mathcal{V}} p_{\mathcal{G}_{\mathcal{V}}}(x_d|\boldsymbol{\phi}_{[d,k]})}{\sum_{k'} n_{[\mathcal{V},k']}} \right). \tag{6}$$

The CrossCat prior for parameters and generative process for data $\mathcal{D} = \{\boldsymbol{x}_i\}_{i=1}^N$ is summarized below.

CrossCat-Prior

$\alpha' \sim \text{Gamma}(1, 1)$

$\boldsymbol{v} = (v_1, \ldots, v_D) \sim \text{CRP}(\cdot|\alpha')$

$\mathcal{V}_k \leftarrow \{i \in [D] : v_i = k\} \qquad k \in \text{Unique}(\boldsymbol{v})$

$\left\{x_{[i,\mathcal{V}_k]}\right\}_{i=1}^N \sim \text{DpMixture-Prior} \quad k \in \text{Unique}(\boldsymbol{v})$

We refer to [18, 22] for algorithms for posterior inference in CrossCat, and assume we have a set of approximate samples $\left\{\hat{\mathcal{G}}_i\right\}_{i=1}^H$ of all latent CrossCat parameters from the posterior $\pi(\cdot|\mathcal{D})$.

3.1.1 Optimizing a CMI Query

The following lemma shows how CrossCat induces sparsity given a multivariate CMI query.

Lemma 1 *Let $\mathcal{G}$ be a posterior sample from CrossCat, whose full joint distribution is given by (5) and (6). Then, for all $\mathcal{A}, \mathcal{B}, \mathcal{C} \subseteq [D]$.*

$$\mathcal{I}_{\mathcal{G}}\left(\boldsymbol{x}_{\mathcal{A}} : \boldsymbol{x}_{\mathcal{B}} \mid \hat{\boldsymbol{x}}_{\mathcal{C}}\right) = \sum_{\mathcal{V} \in \boldsymbol{\gamma}} \mathcal{I}_{\mathcal{G}_{\mathcal{V}}}(\boldsymbol{x}_{\mathcal{A}\cap\mathcal{V}} : \boldsymbol{x}_{\mathcal{B}\cap\mathcal{V}} \mid \hat{\boldsymbol{x}}_{\mathcal{C}\cap\mathcal{V}}),$$

where $\mathcal{I}_{\mathcal{G}_{\mathcal{V}}}(\boldsymbol{x}_{\mathcal{A}\cap\mathcal{V}} : \varnothing|\hat{\boldsymbol{x}}_{\mathcal{C}\cap\mathcal{V}}) \equiv 0$.

Proof

$$\begin{aligned}
\mathcal{I}_{\mathcal{G}}\left(\boldsymbol{x}_{\mathcal{A}}{:}\boldsymbol{x}_{\mathcal{B}}|\hat{\boldsymbol{x}}_{\mathcal{C}}\right) &= \mathbb{E}\left[\log\left(\frac{p_{\mathcal{G}}(\boldsymbol{x}_{\mathcal{A}}{:}\boldsymbol{x}_{\mathcal{B}}|\hat{\boldsymbol{x}}_{\mathcal{C}})}{p_{\mathcal{G}}(\boldsymbol{x}_{\mathcal{A}}|\hat{\boldsymbol{x}}_{\mathcal{C}})p_{\mathcal{G}}(\hat{\boldsymbol{x}}_{\mathcal{B}}|\hat{\boldsymbol{x}}_{\mathcal{C}})}\right)\right] \\
&= \mathbb{E}\left[\log\left(\prod_{\mathcal{V}\in\boldsymbol{\gamma}}\left(\frac{p_{\mathcal{G}_{\mathcal{V}}}(\hat{\boldsymbol{x}}_{\mathcal{A}\cap\mathcal{V}}, \hat{\boldsymbol{x}}_{\mathcal{B}\cap\mathcal{V}}|\hat{\boldsymbol{x}}_{\mathcal{C}\cap\mathcal{V}})}{p_{\mathcal{G}_{\mathcal{V}}}(\hat{\boldsymbol{x}}_{\mathcal{A}\cap\mathcal{V}}|\hat{\boldsymbol{x}}_{\mathcal{C}\cap\mathcal{V}})p_{\mathcal{G}_{\mathcal{V}}}(\hat{\boldsymbol{x}}_{\mathcal{B}\cap\mathcal{V}}|\hat{\boldsymbol{x}}_{\mathcal{C}\cap\mathcal{V}})}\right)\right)\right] \\
&= \mathbb{E}\left[\sum_{\mathcal{V}\in\boldsymbol{\gamma}}\left(\log\left(\frac{p_{\mathcal{G}_{\mathcal{V}}}(\hat{\boldsymbol{x}}_{\mathcal{A}\cap\mathcal{V}}, \hat{\boldsymbol{x}}_{\mathcal{B}\cap\mathcal{V}}|\hat{\boldsymbol{x}}_{\mathcal{C}\cap\mathcal{V}})}{p_{\mathcal{G}_{\mathcal{V}}}(\hat{\boldsymbol{x}}_{\mathcal{A}\cap\mathcal{V}}|\hat{\boldsymbol{x}}_{\mathcal{C}\cap\mathcal{V}})p_{\mathcal{G}_{\mathcal{V}}}(\hat{\boldsymbol{x}}_{\mathcal{B}\cap\mathcal{V}}|\hat{\boldsymbol{x}}_{\mathcal{C}\cap\mathcal{V}})}\right)\right)\right] \\
&= \sum_{\mathcal{V}\in\boldsymbol{\gamma}}\left(\mathbb{E}\left[\log\left(\frac{p_{\mathcal{G}_{\mathcal{V}}}(\hat{\boldsymbol{x}}_{\mathcal{A}\cap\mathcal{V}}, \hat{\boldsymbol{x}}_{\mathcal{B}\cap\mathcal{V}}|\hat{\boldsymbol{x}}_{\mathcal{C}\cap\mathcal{V}})}{p_{\mathcal{G}_{\mathcal{V}}}(\hat{\boldsymbol{x}}_{\mathcal{A}\cap\mathcal{V}}|\hat{\boldsymbol{x}}_{\mathcal{C}\cap\mathcal{V}})p_{\mathcal{G}_{\mathcal{V}}}(\hat{\boldsymbol{x}}_{\mathcal{B}\cap\mathcal{V}}|\hat{\boldsymbol{x}}_{\mathcal{C}\cap\mathcal{V}})}\right)\right)\right] \\
&= \sum_{\mathcal{V}\in\boldsymbol{\gamma}}\mathcal{I}_{\mathcal{G}_{\mathcal{V}}}\left(\boldsymbol{x}_{\mathcal{A}\cap\mathcal{V}}{:}\boldsymbol{x}_{\mathcal{B}\cap\mathcal{V}}|\hat{\boldsymbol{x}}_{\mathcal{C}\cap\mathcal{V}}\right).
\end{aligned}$$

An immediate consequence of Lemma 1 is that structure discovery in CrossCat allows us to optimize Monte Carlo estimation of $\mathcal{I}_{\mathcal{G}}(\boldsymbol{x}_{\mathcal{A}}{:}\boldsymbol{x}_{\mathcal{B}}|\boldsymbol{x}_{\mathcal{C}}{=}\hat{\boldsymbol{x}}_{\mathcal{C}})$ by ignoring all target and condition variables which are not in the same block $\mathcal{V}$, as shown in Algorithm 4a.

Algorithm 4a CrossCat-Cmi

Require: CrossCat $\mathcal{G}$; query $\mathcal{A}$, $\mathcal{B}$; condition $\hat{\boldsymbol{x}}_{\mathcal{C}}$; acc. T
Ensure: Monte Carlo estimate of $\mathcal{I}_{\mathcal{G}}(\boldsymbol{x}_{\mathcal{A}}{:}\boldsymbol{x}_{\mathcal{B}}|\boldsymbol{x}_{\mathcal{C}}{=}\hat{\boldsymbol{x}}_{\mathcal{C}})$

```
for V ∈ γ do
  if A ∩ V and B ∩ V then
    i_V ← Gpm-Cmi(G_V, A ∩ V, B ∩ V, x̂_{C∩V}, T)
  else
    i_V ← 0
return Σ_{V∈γ} i_V
```

3.1.2 Upper bounding dependence probability

In exploratory data analysis, we are often interested in detecting pairwise predictive relationships between variables (x_i, x_j). Using the formalism from Eq (3), the quantity of interest can be phrased as the probability that their MI is non-zero: $\mathbb{P}\left[\mathcal{I}_{\mathcal{G}}(x_i : x_j) > 0\right]$. This quantity can be upper-bounded by the posterior probability that x_i and x_j have the same assignments v_i and v_j in the CrossCat variable partition $\boldsymbol{\gamma}$:

$$\begin{aligned}
&\mathbb{P}\left[\mathcal{I}_{\mathcal{G}}(x_i : x_j) > 0\right] \\
&= \mathbb{P}\left[\mathcal{I}_{\mathcal{G}}(x_i : x_j) > 0 \mid \{\mathcal{G} : v_i = v_j\}\right]\mathbb{P}\left[\{\mathcal{G} : v_i = v_j\}\right] \\
&\quad + \mathbb{P}\left[\mathcal{I}_{\mathcal{G}}(x_i : x_j) > 0 \mid \{\mathcal{G} : v_i \neq v_j\}\right]\mathbb{P}\left[\{\mathcal{G} : v_i \neq v_j\}\right] \\
&= \mathbb{P}\left[\mathcal{I}_{\mathcal{G}}(x_i : x_j) > 0 \mid \{\mathcal{G} : v_i = v_j\}\right]\mathbb{P}\left[\{\mathcal{G} : v_i = v_j\}\right] \\
&< \mathbb{P}\left[\{\mathcal{G} : v_i = v_j\}\right] \approx \frac{1}{H}\sum_{h=1}^{H}\mathbb{I}[\hat{\mathcal{G}}_h : \hat{v}_{[h,i]} = \hat{v}_{[h,j]}]. \quad (7)
\end{aligned}$$

The addend in line 3 is zero from Lemma 1. The summand in (7) can be computed in $O(1)$ for CrossCat sample $\hat{\mathcal{G}}_h$. When dependencies among $(x_1, \ldots, x_D)$ are sparse such that many pairs (x_i, x_j) have MI upper bounded by 0, or small tolerance ϵ, (see Figure 5c), the number of invocations of Algorithm 4a required to tighten the pairwise MI bounds is $\ll O(D^2)$.

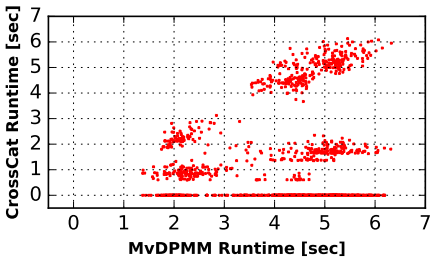

Figure 3: Comparing the runtime of CrossCat-Cmi (Alg 4a) and Gpm-Cmi (Alg 2a) (for the vanilla DpMixture), on 1000 randomly generated CMI queries from an 8-dimensional dataset with T=100. The horizontal line at 0 corresponds to CrossCat detecting structural independence between the query variables, bypassing Monte Carlo estimation completely. The dots below the diagonal correspond to CrossCat optimizing the Monte Carlo estimator, by ignoring any constraint variables which are structurally independent of the target variables. The dots along the diagonal correspond to cases where CrossCat cannot simplify the query, requiring full Monte Carlo estimation and resulting in comparable runtime to DpMixture. These three cases correspond to the three illustration of posterior CrossCat structures and annotated query analysis blocks shown in Figure 1.

mech	vectors	algebra	analysis	stats
77	82	67	67	81
23	38	36	48	15
63	78	80	70	81
55	72	63	70	68
...	...	...	...	...

(a) Student exam scores

	M	V	G	L	S
M	1.00	0.33	0.23	0.00	0.03
V	0.33	1.00	0.28	0.08	0.02
G	0.23	0.28	1.00	0.43	0.36
L	0.00	0.08	0.43	1.00	0.26
S	0.02	0.02	0.36	0.26	1.00

(b) Partial correlation matrix; red entries indicate statistically significant conditional independence.

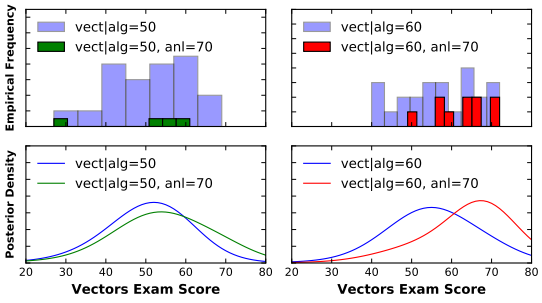

(c) Empirical frequencies from the dataset and posterior predictive distributions from CrossCat.

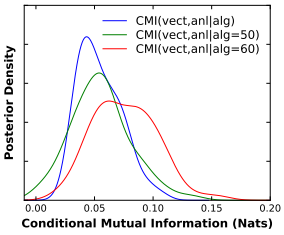

(d) Posterior distributions of CMI(vectors, analysis) given various conditions of algebra.

Figure 4: Using posterior CMI distributions to discover conditional and context-specific predictive relationships in the mathematics marks dataset. **(b)** The red entries in the partial correlation matrix indicates that vectors is conditionally independent of analysis given algebra. **(c)** The left panel shows that when algebra = 50, conditioning on analysis = 70 appears to have little effect on the prediction for vectors. The right panel shows that when algebra = 60, however, conditioning on analysis = 70 results in a big shift of the posterior mean of vectors from 52 to just under 70. This shift is consistent with the top right histogram, where knowing that analysis = 70 eliminates all empirical vector scores in the heavy left tail. **(d)** The distribution of CMI of vectors and analysis when analysis = 60 (red curve) places more density at higher values of information than when algebra = 50 (green curve), explaining the shift in (c). The blue curve shows the CMI marginalizing over algebra.

4 Applications

This section illustrates the efficacy of the proposed approach on (i) a real-world dataset of mathematics exam scores and (ii) a sparse database containing ∼320 indicators of global macroeconomic development drawn from an ongoing collaboration with the Bill & Melinda Gates Foundation.

4.1 Mathematics exam scores

The "mathematics marks" dataset contains 5 examination scores for 88 students: mechanics, vectors, algebra, analysis, and statistics. It is a canonical example in the literature of a real-valued, multivariate dataset with interesting conditional independence structure [19, 32, 6]. Modeling the variables as jointly Gaussian and inspecting the partial correlation matrix (Figure 4b) indicates that (mechanics and vectors) are conditionally independent of (analysis and statistics), given algebra. The partial correlation matrix implies that when predicting the vectors score for a student whose algebra score is known, further conditioning on the analysis score provides no additional information.

Figure 4 explores this proposed conditional independence between vectors and analysis by studying (i) the empirical frequencies from the dataset, (ii) the posterior predictive distributions from CrossCat, and (iii) the posterior distribution of CMI given two different conditions: when algebra = 50, and when algebra = 60. The posterior CMI curves in Figure 4d detects a "context-specific" dependence, where analysis is more predictive of vectors when algebra = 60 (red curve), than when algebra = 50 (green curve). The CMI of vectors and analysis is weakest when marginalizing over all values of algebra (blue curve in 4d), in agreement with their weak partial correlation. The figure caption contains further commentary.

4.2 Macroeconomic indicators of global poverty, education, and health

The Gapminder data set is an extensive longitudinal dataset of over 300 global developmental indicators for 200 countries spanned over 5 centuries [26]. These include variables from a broad set of categories such as education, health, trade, poverty, population growth, and mortality rates. Our experiments are restricted to a cross-sectional slice of the data from 2002. Figure 5a shows the pairwise R^2 correlation values between all variables; each row and column in the heatmap is an indicator in the dataset, and the color of a cell is the raw value of R^2 (between 0 and 1). Figure 5b shows pairwise binary hypothesis tests of independence using the HSIC [12], which detects a dense set of dependencies including many spurious relationships (see supplement). For these two methods, statistically insignificant relationships ($\alpha = 0.05$ with Bonferroni correction for multiple testing) have been converted to 0. Figure 5c shows the pairwise posterior probability that the MI of two variables is greater than zero (also a value between 0 and 1). These entries are estimated using Eq (7) with H=100 samples of CrossCat.

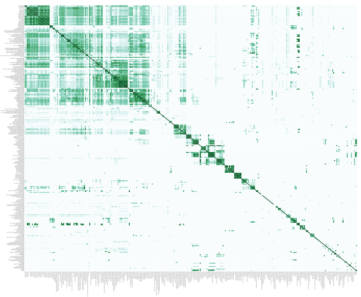

(a) R^2 Correlation Value

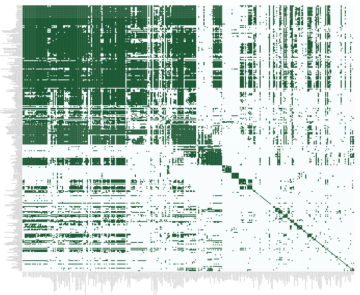

(b) HSIC [12] Independence Test

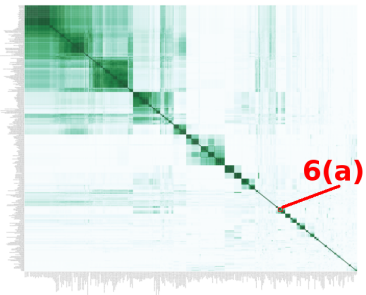

(c) $\mathbb{P}[\mathcal{I}_{\mathcal{G}}(x_i{:}x_j) > 0]$, Eq (7)

(d) Pairwise heatmaps of all 320 variables in the Gapminder dataset using three dependency detection techniques. Darker cells indicate a detected dependence between the two variables.

variable A	variable B	Pr[I>0]	R^2 ($p \ll 10^{-6}$)
personal computers	earthquake affected	0.015625	0.974445
road traffic total deaths	people living w/ hiv	0.265625	0.858260
natural gas reserves	15-25 yrs sex ratio	0.031250	0.951471
forest products per ha	earthquake killed	0.046875	0.936342
flood affected	population 40-59 yrs	0.140625	0.882729

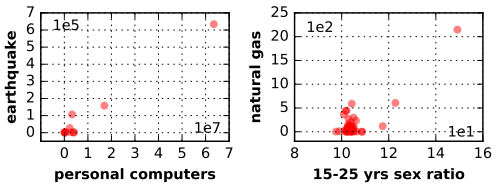

(e) Spurious relationships which are correlated under R^2, but probably independent according to posterior MI.

variable A	variable B	Pr[I>0]	R^2 ($p \ll 10^{-6}$)
inflation	trade balance (% gdp)	0.859375	0.114470
DOTS detection rate	DOTS coverage	0.812500	0.218241
forest area (sq. km)	forest products (usd)	0.828125	0.206145
long term unemp. rate	total 15-24 unemp.	0.968750	NaN
male family workers	female self-employed	0.921875	NaN

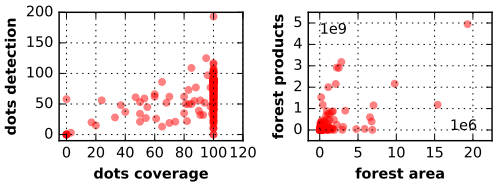

(f) Common-sense relationships probably dependent according to posterior MI, but weakly correlated under R^2.

Figure 5: Comparing dependences between variables in the Gapminder dataset, as detected by R^2, HSIC (with Bonferroni correction for multiple testing), and posterior distribution over mutual information in CrossCat.

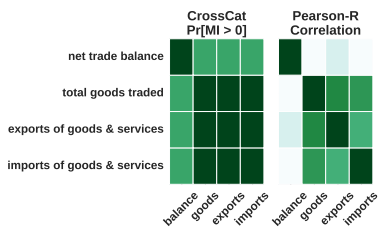

(g) Block of "trade" variables detected as probably dependent.

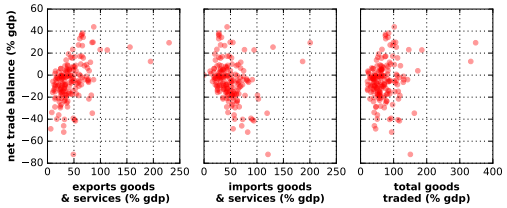

(h) Scatter plots show weak linear correlations and heteroskedastic noise for `net balance`.

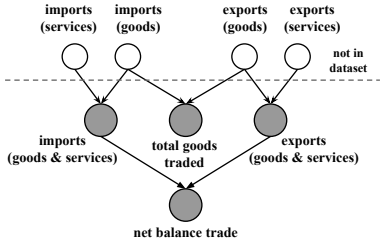

(i) Ground truth causal structure of variables in the "trade" block.

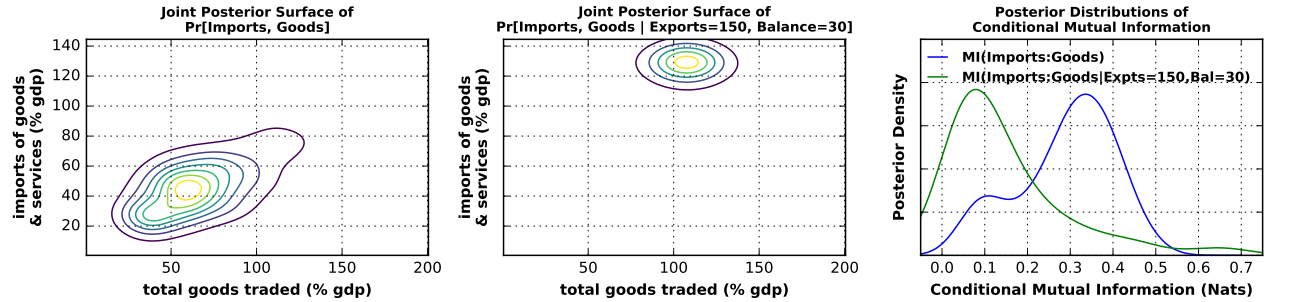

(j) Left plot shows the joint posterior density of imports and goods; the high marginal coupling is due to their common parent in (i). Center plot shows the same distribution conditioned on exports=150 and balance=30; imports now centers around its noiseless value of 120, and is decoupled from goods. Right plot shows the CMI for these two distributions.

Figure 6: CMI discovers existence and confirms strength of predictive relationships between "trade" variables.

They are computed without Monte Carlo integration by inspecting only the variable partition discovered by each CrossCat sample. It is important to note that the metric Pr[MI(A:B)>0] in Figure 5c only indicates the possible *existence* of a predictive relationship between variables A and B, and makes no statement about the strength or directionality of the relationship.

It is instructive to compare the dependencies detected by R^2 and CrossCat-Cmi. Table 5e shows pairs of variables that are spuriously reported as dependent according to correlation; scatter plots reveal they are either (i) are sparsely observed or (ii) exhibit high correlation due to large outliers. Table 5f shows common-sense relationships between pairs of variables that CrossCat-CMI detects but R^2 does not; scatter plots reveal they are either (i) non-linearly related, (ii) roughly linear with heteroskedastic noise, or (iii) pairwise independent but dependent given a third variable. Recall that CrossCat is a product of DpMixtures; practically meaningful conditions for weak and strong consistency of Dirichlet location-scale mixtures have been established by [10, 31]. This supports the intuition that CrossCat can detect a broad class of predictive relationships that simpler parametric models miss.

Figure 6 focuses on a group of four "trade"-related variables in the Gapminder dataset detected as probably dependent: "net trade balance", "total goods traded", "exports of goods and services", and "imports of goods and services". R^2 fails to detect a statistically significant dependence between "net trade balance" and the other variables, due to the high noise shown in the pairwise scatter plots (Figure 5h). From economics, these four variables are related by the graphical model in Figure 5i, where the value of a node is a noisy addition or subtraction of the values of its parents. Figure 5j illustrates that CrossCat recovers predictive relationships between these variables: conditioning on "exports"=150 and "balance"=30 (a low probability event according to the left subplot) centers the posterior predictive distribution of "imports" around 120, and decouples it from "total goods". The posterior CMI curves of "imports" and "total goods", with and without the conditions on "exports" and "balance", formalize this decoupling (right subplot of Figure 5j).

5 Related Work

There is broad acknowledgment that new techniques for dependency detection beyond linear correlation are required. Frequentist approaches for conditional independence testing include the use of kernel methods [1, 9, 33, 28], copula functions [2, 24, 16], and characteristic functions [30]; a comparison can be found in [16, Table 1]. Several of these non-parametric methods aim to capture non-linear and multivariate predictive relationships. Unlike these methods, however, our approach uses a representation of dependence in terms of conditional mutual information. Also, the proposed method is not embedded in a frequentist decision-theoretic framework, since our quantity of interest is a posterior distribution over CMI rather than a binary hypothesis test where the null is CMI = 0.

In the Bayesian literature, the most similar approach to ours was proposed independently in recent work by [15]. The method likewise computes a posterior distribution over CMI, estimating the joint density using an encompassing non-parametric Bayesian prior. However, the differences are significant. First, the Monte Carlo CMI estimator in [15] is based on resampling the empirical data, and its consistency is established under regularity conditions. However, real-world databases may be too sparse for resampling to yield good estimates, especially for queries given unlikely constraints. Gpm-Cmi instead uses simulations from a GPM, and CrossCat-Cmi in particular uses samples from the predictive distribution in a non-parametric model. Second, the non-parametric prior used by [15] is a standard Dirichlet process mixture model, whereas this paper proposes a sparsity-inducing CrossCat prior which permits (i) optimized computations for upper bounds of posterior probabilities and (ii) simplifying complex CMI queries with multivariate conditions.

6 Discussion

This paper has shown it is possible to detect predictive relationships between variables by integrating probabilistic programming, information theory, and non-parametric Bayes. Users can specify a broad class of conditional mutual information queries using a simple SQL-like language (see supplement), which are answered using a scalable pipeline based on approximate Bayesian inference. The underlying approach applies to arbitrary GPMs, including parametric models and other classes of probabilistic programs [27], but exploits the sparsity of CrossCat model structures to improve scalability for exploratory analysis.

As scientific research and business operations become more data-driven, there is a growing need for statistically rigorous ways to characterize predictive relationships. This paper has focused on detection and preliminary assessment of the probable strength of a relationship in information-theoretic terms. Confirmatory analysis and descriptive summarization are left for future work, and will require new ideas, including insight into the robustness of joint density estimation when random sampling assumptions are violated. Despite these limitations, we believe the proposed method represents a significant step towards automatic methods for identifying relationships from empirical data.

References

[1] Francis R Bach and Michael I Jordan. Kernel independent component analysis. *Journal of machine learning research*, 3(Jul):1–48, 2002.

[2] Taoufik Bouezmarni, Jeroen VK Rombouts, and Abderrahim Taamouti. Nonparametric copula-based test for conditional independence with applications to granger causality. *Journal of Business & Economic Statistics*, 30(2):275–287, 2012.

[3] Craig Boutilier, Nir Friedman, Moises Goldszmidt, and Daphne Koller. Context-specific independence in bayesian networks. In *Proceedings of the Twelfth international conference on Uncertainty in artificial intelligence*, pages 115–123. Morgan Kaufmann Publishers Inc., 1996.

[4] N Council. *Frontiers in massive data analysis*. The National Academies Press Washington, DC, 2013.

[5] T.M. Cover and J.A. Thomas. *Elements of Information Theory*. Wiley, 2012. ISBN 9781118585771.

[6] David Edwards. *Introduction to graphical modelling*. Springer Science & Business Media, 2012.

[7] Michael D Escobar and Mike West. Bayesian density estimation and inference using mixtures. *Journal of the american statistical association*, 90 (430):577–588, 1995.

[8] Daniel Fink. A compendium of conjugate priors. 1997.

[9] Kenji Fukumizu, Arthur Gretton, Xiaohai Sun, and Bernhard Schölkopf. Kernel measures of conditional dependence. In *NIPS*, volume 20, pages 489–496, 2007.

[10] Subhashis Ghosal, Jayanta K Ghosh, RV Ramamoorthi, et al. Posterior consistency of dirichlet mixtures in density estimation. *The Annals of Statistics*, 27(1):143–158, 1999.

[11] Dilan Görür and Carl Edward Rasmussen. Dirichlet process gaussian mixture models: Choice of the base distribution. *Journal of Computer Science and Technology*, 25(4):653–664, 2010.

[12] Arthur Gretton, Olivier Bousquet, Alex Smola, and Bernhard Schölkopf. Measuring statistical dependence with hilbert-schmidt norms. In *International conference on algorithmic learning theory*, pages 63–77. Springer, 2005.

[13] Sonia Jain and Radford M Neal. A split-merge markov chain monte carlo procedure for the dirichlet process mixture model. *Journal of Computational and Graphical Statistics*, 2012.

[14] Alexander Kraskov, Harald Stögbauer, and Peter Grassberger. Estimating mutual information. *Physical review E*, 69(6):066138, 2004.

[15] Tsuyoshi Kunihama and David B Dunson. Nonparametric bayes inference on conditional independence. *Biometrika*, 103(1):35–47, 2016.

[16] David Lopez-Paz, Philipp Hennig, and Bernhard Schölkopf. The randomized dependence coefficient. In *Advances in Neural Information Processing Systems*, pages 1–9, 2013.

[17] V. Mansinghka, R. Tibbetts, J. Baxter, P. Shafto, and B. Eaves. Bayesdb: A probabilistic programming system for querying the probable implications of data. *arXiv preprint arXiv:1512.05006*, 2015.

[18] Vikash Mansinghka, Patrick Shafto, Eric Jonas, Cap Petschulat, Max Gasner, and Joshua B. Tenenbaum. Crosscat: A fully bayesian nonparametric method for analyzing heterogeneous, high dimensional data. *Journal of Machine Learning Research*, 17(138):1–49, 2016.

[19] Kantilal Varichand Mardia, John T Kent, and John M Bibby. *Multivariate analysis*. Academic press, 1980.

[20] Rudy Moddemeijer. On estimation of entropy and mutual information of continuous distributions. *Signal processing*, 16(3):233–248, 1989.

[21] Young-Il Moon, Balaji Rajagopalan, and Upmanu Lall. Estimation of mutual information using kernel density estimators. *Physical Review E*, 52(3): 2318, 1995.

[22] Fritz Obermeyer, Jonathan Glidden, and Eric Jonas. Scaling nonparametric bayesian inference via subsample-annealing. In *Proceedings of the Seventeenth International Conference on Artificial Intelligence and Statistics, AISTATS 2014, Reykjavik, Iceland, April 22-25, 2014*, pages 696–705, 2014.

[23] Liam Paninski. Estimation of entropy and mutual information. *Neural computation*, 15(6):1191–1253, 2003.

[24] Barnabás Póczos, Zoubin Ghahramani, and Jeff Schneider. Copula-based kernel dependency measures. *arXiv preprint arXiv:1206.4682*, 2012.

[25] C. Robert and G. Casella. *Monte Carlo Statistical Methods*. Springer Texts in Statistics. Springer New York, 2005. ISBN 9780387212395.

[26] Hans Rosling. Gapminder: World. 2008.

[27] Feras Saad and Vikash Mansinghka. Probabilistic data analysis with probabilistic programming. *CoRR*, abs/1608.05347, 2016.

[28] Dino Sejdinovic, Arthur Gretton, and Wicher Bergsma. A kernel test for three-variable interactions. In *Advances in Neural Information Processing Systems*, pages 1124–1132, 2013.

[29] Ross D Shachter. Bayes-ball: Rational pastime (for determining irrelevance and requisite information in belief networks and influence diagrams). In *Proceedings of the Fourteenth conference on Uncertainty in artificial intelligence*, pages 480–487. Morgan Kaufmann Publishers Inc., 1998.

[30] Liangjun Su and Halbert White. A consistent characteristic function-based test for conditional independence. *Journal of Econometrics*, 141(2): 807 – 834, 2007. ISSN 0304-4076.

[31] Surya T Tokdar. Posterior consistency of dirichlet location-scale mixture of normals in density estimation and regression. *Sankhyā: The Indian Journal of Statistics*, pages 90–110, 2006.

[32] J. Whittaker. *Graphical models in applied multivariate statistics*. Wiley series in probability and mathematical statistics: Probability and mathematical statistics. Wiley, 1990. ISBN 9780471917502.

[33] Kun Zhang, Jonas Peters, and Dominik Janzing. Kernel-based conditional independence test and application in causal discovery. In *In Uncertainty in Artificial Intelligence*, pages 804–813, 2011.

A Expressing conditional mutual information queries in the Bayesian Query Language

In this section, we provide technical details outlining the integration of the workflow in Figure 1 into BayesDB [17], a probabilistic programming platform for applied data analysis. We also outline new syntaxes in the Bayesian Query Language, an SQL-like query language which allows users to ask probabilistic questions about the data, to support posterior CMI simulations.

The data table in block 1 of Figure 1 is referred to as a *base table* $\mathcal{D}$ in BayesDB, which contains raw data records stored in tabular format. A *population* $\mathcal{P}$ for $\mathcal{D}$ provides each variable (column) in the dataset with a a statistical data type, such as nominal, cyclic, magnitude, counts, and real [27, Table 2]. The statistical data type of a variable indicates the set of allowable values that its realizations can take (such as the integers for the counts data type, or the positive reals for the magnitude data type), as well as the qualitative invariants they satisfy (such as permutation invariance for a nominal data type). Finally, a population $\mathcal{P}$ is associated with a *metamodel* $\mathcal{M}$, which is a container for the ensemble of probabilistic models used to answer BQL queries about $\mathcal{P}$. A complete description of this architecture is provided in [27, Section 5]. Formally, the metamodel $\mathcal{M} = \{(\mathcal{G}_i, w_i)\}_{i=1}^{K}$ is a weighted ensemble of GPMs, whose minimal computational interface is shown in Listing 1.

When each $\mathcal{G}_i$ is a posterior sample (given data in the base table $\mathcal{D}$) from a Bayesian model, and the weights $w_1 = \cdots = w_K$ are equal, then functionals of $\mathcal{G}_i$ (such as the conditional mutual information (2)) are random variables whose (approximate) posterior distribution arises from evalauting the functional at each GPM.

To expose the posterior distribution of CMI to BayesDB, we extend the Bayesian Query Language with the following grammar

```
query       ::=
  SIMULATE CONDITIONAL MUTUAL INFORMATION OF
  target WITH target
  given
  USING <T> SAMPLES
  FROM MODELS OF <population>.

target            ::= <var-name>.
target            ::= target, <var-name>.

given             ::= GIVEN constraints | \0.
constraints       ::= var-constraint.
constraints       ::= constraints, var-constraint.

var-constraint    ::= e-constraint | m-constraint.
e-constraint      ::= <var-name>=<var-val>.
m-constraint      ::= <var-name>.
```

Red symbols denote non-terminals, upper case symbols in green denote terminals, lower-case symbols between angle brackets denote literals, and `\0` is the empty string. The grammar extends the original specification of BQL `SIMULATE` from [17, Section 3.2.1], which only allowed users to directly simulate variables in the population, rather than arbitrary derived random variables defined as functions of the population GPMs. To illustrate the SQL analogy of the grammar, we may conceptualize the expression `FROM MODELS OF <population>` as specifying a one-column table, where row i is $\mathcal{G}_i$. The `CONDITIONAL MUTUAL INFORMATION` is thus a one-row function of this table, parameterized by the user-specified `target` and `given` variables.

The grammar supports *context-specific* CMI queries (2), where conditioning variables have realizations,

```
SIMULATE CONDITIONAL MUTUAL INFORMATION OF
  (x1, x2) WITH x3 GIVEN (x4 = 14, x5 = 1)
  USING 100 SAMPLES
  FROM MODELS OF pop;
```

marginalization CMI queries (4), where conditioning variables are marginalized out,

```
SIMULATE CONDITIONAL MUTUAL INFORMATION OF
  (x1, x2) WITH x3 GIVEN (x6, x7)
  USING 100 SAMPLES
  FROM MODELS OF pop;
```

and an arbitrary combination of the two,

```
SIMULATE CONDITIONAL MUTUAL INFORMATION OF
  (x1, x2) WITH x3 GIVEN (x4 = 14, x5 = 1, x6, x7)
  USING 100 SAMPLES
  FROM pop;
```

All of these queries return samples from the posterior distribution over CMI, (shown as rug ticks in block 8 of 1). The number of returned samples is K, or the total number of GPMs in metamodel $\mathcal{M}$.

Finally, we note that the `given` expression only supports equality constraints and marginalization constraints. Morere elaborate probabilistic conditioning clauses, such as `(GIVEN <var-name> BETWEEN <low> AND <high>)` are left to future work. Approaches to computationally answering such CMI queries with set-containment constraints might involve rejection sampling, or extending the GPM interface to assessing cumulative distribution functions when possible.

B Experimental Methods For Dependence Detection Baselines

In this section we outline the methodology used to produce the pairwise R^2 and HSIC heatmaps shown in Figures 5a and 5b. To detect the strength of linear correlation (for R^2) and perform a marginal independence test (for HSIC) given variables x_i and x_j in the Gapminder dataset, all records in which at least one of these two variables is missing were dropped. If the total number of remaining observations was less than three, the null hypothesis of independence was not rejected due to degeneracy of these methods at very small sample sizes. Hypothesis tests were performed at the $\alpha = 0.05$ significance level. To account for multiple testing (a total of $\binom{320}{2} = 51040$), a standard Bonferroni correction was applied to ensure a family-wise error rate of at most α.

We used an open source MATLAB implementation for HSIC.[1] 1000 permutations were used to approximate the null distribution, and kernel sizes were determined using median distances from the dataset. From Figure 5b, HSIC detects a large number of statistically significant dependencies. Table 1 reports a collection of spurious relationships between variables for which (i) HSIC rejects the null hypothesis of independence, and (ii) the upper bound on the posterior probability of dependence Eq (7) is less than 0.15.

variable A	variable B
body mass index (men)	privately owned forest (%)
energy use (per capita)	50+ yrs sex ratio
co2 emission per person (tons)	total sex ratio
homicide (15-29)	inflation (annual %)
children out of primary school	male above 60 (%)
income share (fourth 20%)	suicide (45-59)
personal computers	arms exports (us$)
agricultural land (% land area)	total population female
billionaires (per 1M)	mobile subscription (per 100)
residential elec. consumption	smear-positive detection (%)
coal consumption (per capita)	age at 1st marriage (women)
underweight children	murder (per 100,000), age adj.
RTI 60+	oil consumption per capita
energy use (per capita)	number of child deaths
female 0-4 years (%)	15+ literacy rate (%)
broadband subscribers (%)	dots new case detection (%)
dead kids per woman	coal consumption per person
crude oil prod. (per capita)	TB incidence (%)
dependency ratio	people living with hiv

Table 1: Spurious relationships detected as dependent by HSIC, but probably independent according to posterior mutual information.

[1] Function `hsicTestBoot` from `http://www.gatsby.ucl.ac.uk/~gretton/indepTestFiles/indep.htm`.